Terms used **initiate recovery lock NFS** Found **4** of **72 searched out of 72.**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 4 of 4

Relevance scale

1 Recovery in the Calypso file system
Murthy Devarakonda, Bill Kish, Ajay Mohindra
August 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 3
Publisher: ACM Press
Full text available: pdf(318.88 KB) Additional Information: full citation, abstract, references, citings, index terms, review

This article presents the deign and implementation of the recovery scheme in Calypso. Calypso is a cluster-optimized, distributed file system for UNIX clusters. As in Sprite and AFS, Calypso servers are stateful and scale well to a large number of clients. The recovery scheme in Calypso is nondisruptive, meaning that open files remain open, client modified data are saved, and in-flight operations are properly handled across server recover. The scheme uses distributed state amount the client ...

Keywords: Calypso, cluster systems, distributed state, state reconstruction

2 Frangipani: a scalable distributed file system
Chandramohan A. Thekkath, Timothy Mann, Edward K. Lee
October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5
Publisher: ACM Press
Full text available: pdf(2.20 MB) Additional Information: full citation, references, citings, index terms

3 Interposed request routing for scalable network storage
February 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 1
Publisher: ACM Press
Full text available: pdf(363.12 KB) Additional Information: full citation, abstract, references, citings, index terms, review

This paper explores interposed request routing in Slice, a new storage system architecture for high-speed networks incorporating network-attached block storage. Slice interposes a request switching filter---called a μ*proxy*---along each client's network path to the storage service (e.g., in a network adapter or switch). The μproxy intercepts request traffic and distributes it across a server ensemble. We propose request routing schemes for I/O and file service traffic, and explore th ...

Keywords: Content switch, file server, network file system, network storage, request redirection, service virtualization

4 Serverless network file systems

Thomas E. Anderson, Michael D. Dahlin, Jeanna M. Neefe, David A. Patterson, Drew S. Roselli, Randolph Y. Wang

February 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 1

Publisher: ACM Press

We propose a new paradigm for network file system design: serverless network file systems. While traditional network file systems rely on a central server machine, a serverless system utilizes workstations cooperating as peers to provide all file system services. Any machine in the system can store, cache, or control any block of data. Our approach uses this location independence, in combination with fast local area networks, to provide better performance and scalability th ...

Keywords: RAID, log cleaning, log structured, log-based striping, logging, redundant data storage, scalable performance

Results 1 - 4 of 4

Terms used **network file system lock recover client list** Found **3** of **61** searched out of **1,643.**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 3 of 3

Relevance scale

1 **Distributed file systems: concepts and examples**

Eliezer Levy, Abraham Silberschatz
December 1990 **ACM Computing Surveys (CSUR)**, Volume 22 Issue 4
Publisher: ACM Press

Full text available: pdf(5.33 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The purpose of a distributed file system (DFS) is to allow users of physically distributed computers to share data and storage resources by using a common file system. A typical configuration for a DFS is a collection of workstations and mainframes connected by a local area network (LAN). A DFS is implemented as part of the operating system of each of the connected computers. This paper establishes a viewpoint that emphasizes the dispersed structure and decentralization of both data and con ...

2 **Recovery in the Calypso file system**

Murthy Devarakonda, Bill Kish, Ajay Mohindra
August 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 3
Publisher: ACM Press

Full text available: pdf(318.88 KB) Additional Information: full citation, abstract, references, citings, index terms, review

This article presents the deign and implementation of the recovery scheme in Calypso. Calypso is a cluster-optimized, distributed file system for UNIX clusters. As in Sprite and AFS, Calypso servers are stateful and scale well to a large number of clients. The recovery scheme in Calypso is nondisruptive, meaning that open files remain open, client modified data are saved, and in-flight operations are properly handled across server recover. The scheme uses distributed state amount the client ...

Keywords: Calypso, cluster systems, distributed state, state reconstruction

3 **High-latency, low-bandwidth windowing in the Jupiter collaboration system**

David A. Nichols, Pavel Curtis, Michael Dixon, John Lamping
December 1995 **Proceedings of the 8th annual ACM symposium on User interface and software technology**
Publisher: ACM Press

Full text available: pdf(1.03 MB) Additional Information: full citation, references, citings, index terms

Keywords: CSCW, UIMS, groupware toolkits, optimistic currency control, window toolkits

Results 1 - 3 of 3

Terms used **lock recovery** Found **14** of **175,083**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 14 of 14

Relevance scale

1 Scalable fault-tolerant distributed shared memory

Florin Sultan, Liviu Iftode, Thu Nguyen
November 2000 **Proceedings of the 2000 ACM/IEEE conference on Supercomputing (CDROM)**
Publisher: IEEE Computer Society
Full text available: pdf(247.40 KB) Additional Information: full citation, abstract, references, citings, index terms
Publisher Site

This paper shows how a state-of-the-art software distributed shared-memory (DSM) protocol can be efficiently extended to tolerate single-node failures. In particular, we extend a home-based lazy release consistency (HLRC) DSM system with independent checkpointing and logging to volatile memory, targeting shared-memory computing on very large LAN-based clusters. In these environments, where global coordination may be expensive, independent checkpointing becomes critical to scalability. Howev ...

2 A history and evaluation of System R

Donald D. Chamberlin, Morton M. Astrahan, Michael W. Blasgen, James N. Gray, W. Frank King, Bruce G. Lindsay, Raymond Lorie, James W. Mehl, Thomas G. Price, Franco Putzolu, Patricia Griffiths Selinger, Mario Schkolnick, Donald R. Slutz, Irving L. Traiger, Bradford W. Wade, Robert A. Yost
October 1981 **Communications of the ACM**, Volume 24 Issue 10
Publisher: ACM Press
Full text available: pdf(1.55 MB) Additional Information: full citation, abstract, references, citings, index terms

System R, an experimental database system, was constructed to demonstrate that the usability advantages of the relational data model can be realized in a system with the complete function and high performance required for everyday production use. This paper describes the three principal phases of the System R project and discusses some of the lessons learned from System R about the design of relational systems and database systems in general.

Keywords: access path selection, authorization, compilation, database management systems, locking, recovery, relational model

3 On deadlocks in interconnection networks

Timothy Mark Pinkston, Sugath Warnakulasuriya
May 1997 **ACM SIGARCH Computer Architecture News , Proceedings of the 24th annual international symposium on Computer architecture ISCA '97**, Volume 25 Issue 2
Publisher: ACM Press

Full text available: pdf(1.87 MB) Additional Information: full citation, abstract, references, citings, index terms

Deadlock avoidance-based and deadlock recovery-based routing algorithms have been proposed in recent years without full understanding of the likelihood and characteristics of actual deadlocks in interconnection networks. This work models the interrelationships between routing freedom, message blocking, correlated resource dependencies and deadlock formation. We empirically show that increasing routing freedom, as achieved by allowing unrestricted routing over multiple virtual channels, makes dea ...

4 Synchronization and recovery in a client-server storage system

E. Panagos, A. Biliris

August 1997 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 6 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(205.25 KB) Additional Information: full citation, abstract, citings, index terms

Client-server object-oriented database management systems differ significantly from traditional centralized systems in terms of their architecture and the applications they target. In this paper, we present the client-server architecture of the EOS storage manager and we describe the concurrency control and recovery mechanisms it employs. EOS offers a semi-optimistic locking scheme based on the multi-granularity two-version two-phase locking protocol. Under this scheme, multiple concurrent reade ...

Keywords: Checkpoint, Client-server architecture, Object management, Concurrency control, Locking, Logging, Recovery, Transaction management

5 The impact of recovery on concurrency control

W. E. Weihl

March 1989 **Proceedings of the eighth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems**

Publisher: ACM Press

Full text available: pdf(1.39 MB) Additional Information: full citation, abstract, references, citings, index terms

It is widely recognized by practitioners that concurrency control and recovery for transaction systems interact in subtle ways. In most theoretical work, however, concurrency control and recovery are treated as separate, largely independent problems. In this paper we investigate the interactions between concurrency control and recovery. We consider two general recovery methods for abstract data types, update-in-place and deferred-update. While each requires operations to conflict if they do ...

6 Atomic data abstractions in a distributed collaborative editing system

Irene Greif, Robert Seliger, William E. Weihl

January 1986 **Proceedings of the 13th ACM SIGACT-SIGPLAN symposium on Principles of programming languages**

Publisher: ACM Press

Full text available: pdf(2.52 MB) Additional Information: full citation, abstract, references, citings

This paper describes our experience implementing CES, a distributed Collaborative Editing System written in Argus, a language that includes facilities for managing long-lived distributed data. Argus provides atomic actions, which simplify the handling of concurrency and failures, and mechanisms for implementing atomic data types, which ensure serializability and recoverability of actions that use them. This paper focuses on the support for atomicity in Argus ...

7 The jitter time-stamp approach for clock recovery of real-time variable bit-rate traffic

Weilian Su, Ian F. Akyildiz

December 2001 **IEEE/ACM Transactions on Networking (TON)**, Volume 9 Issue 6

Publisher: IEEE Press

Full text available: pdf(178.70 KB) Additional Information: full citation, abstract, references, index terms

When multimedia streams arrive at the receiver, their temporal relationships may be distorted due to jitter. Assuming the media stream is packetized, the jitter is then the packet's arrival time deviation from its expected arrival time. There are various ways to reduce jitter, which include synchronization at the application layer, or synchronization at the *asynchronous transfer mode* (ATM) *adaptation layer* (AAL). The new source rate recovery scheme called *jitter time-stamp* (...

Keywords: ATM adaptation layer 2 (AAL2), ATM networks, constant bit-rate (CBR), synchronization, timing recovery, variable bit-rate (VBR)

8 The Recovery Manager of the System R Database Manager

Jim Gray, Paul McJones, Mike Blasgen, Bruce Lindsay, Raymond Lorie, Tom Price, Franco Putzolu, Irving Traiger

June 1981 **ACM Computing Surveys (CSUR)**, Volume 13 Issue 2

Publisher: ACM Press

Full text available: pdf(1.75 MB) Additional Information: full citation, references, citings, index terms

9 Imposing transactional properties on distributed software atchitectures

Apostolos Zarras, Valérie Issarny

September 1998 **Proceedings of the 8th ACM SIGOPS European workshop on Support for composing distributed applications**

Publisher: ACM Press

Full text available: pdf(725.54 KB) Additional Information: full citation, index terms

10 System R: relational approach to database management

M. M. Astrahan, M. W. Blasgen, D. D. Chamberlin, K. P. Eswaran, J. N. Gray, P. P. Griffiths, W. F. King, R. A. Lorie, P. R. McJones, J. W. Mehl, G. R. Putzolu, I. L. Traiger, B. W. Wade, V. Watson

June 1976 **ACM Transactions on Database Systems (TODS)**, Volume 1 Issue 2

Publisher: ACM Press

Full text available: pdf(3.18 MB) Additional Information: full citation, abstract, references, citings, index terms

System R is a database management system which provides a high level relational data interface. The systems provides a high level of data independence by isolating the end user as much as possible from underlying storage structures. The system permits definition of a variety of relational views on common underlying data. Data control features are provided, including authorization, integrity assertions, triggered transactions, a logging and recovery subsystem, and facilities for maintaining ...

Keywords: authorization, data structures, database, index structures, locking, nonprocedural language, recovery, relational model

11 Workflow management based on process model repositories

Volker Gruhn, Monika Schneider

April 1998 **Proceedings of the 20th international conference on Software engineering**

Publisher: IEEE Computer Society

Full text available: pdf(1.12 MB) Publisher Site Additional Information: full citation, references, index terms

12 Concurrency and distribution in object-oriented programming

Jean-Pierre Briot, Rachid Guerraoui, Klaus-Peter Lohr
September 1998 **ACM Computing Surveys (CSUR)**, Volume 30 Issue 3
Publisher: ACM Press

Full text available: pdf(289.34 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper aims at discussing and classifying the various ways in which the object paradigm is used in concurrent and distributed contexts. We distinguish among the library approach, the integrative approach, and the reflective approach. The library approach applies object-oriented concepts, as they are, to structure concurrent and distributed systems through class libraries. The integrative approach consists of merging concepts such as obj ...

Keywords: concurrency, distribution, integration, libraries, message passing, object, reflection

13 Database research at Wisconsin

CORPORATE Univ. of Wisconsin
March 1993 **ACM SIGMOD Record**, Volume 22 Issue 1
Publisher: ACM Press
Full text available: pdf(799.19 KB) Additional Information: full citation, index terms

14 Protocols verified by APPROVER

January 1979 **ACM SIGCOMM Computer Communication Review**, Volume 9 Issue 1
Publisher: ACM Press
Full text available: pdf(142.74 KB) Additional Information: full citation, references

Results 1 - 14 of 14